arXiv:1410.8326v1 [cs.LG] 30 Oct 2014

Towards Learning Object Affordance Priors from Technical Texts

Nicholas H. Kirk[1]

***Abstract*— Everyday activities performed by artificial assistants can potentially be executed naïvely and dangerously given their lack of common sense knowledge. This paper presents conceptual work towards obtaining prior knowledge on the usual modality (passive or active) of any given entity, and their affordance estimates, by extracting high-confidence ability modality semantic relations (*X can Y* relationship) from non-figurative texts, by analyzing co-occurrence of grammatical instances of subjects and verbs, and verbs and objects. The discussion includes an outline of the concept, potential and limitations, and possible feature and learning framework adoption.**

I. Concept

In the domain of autonomous robot control, artificial assistants require to know what actions can be executed on a given set of objects. Such information, defined as *object affordances*, is usually obtained online by reinforcement or active learning during the execution of actions by processing percepts [1]. However, for safe human-robot interaction, we require the robot to have, from initialization, an understanding of *what actions an object can execute*, and *what actions an object can be subject to*. In this scope, we claim human-written technical texts can be an informative source to construct such initial world estimate. Such probability distribution over action-object relationships from natural language text can be performed thanks to the co-occurrence understanding of verb-noun pairs: this analysis is known in computational linguistic literature as the use of distributional information of text to characterize lexical semantics, by considering statistical co-occurrence of neighboring words [2]. However, the majority of current approaches make use of shallow syntactic features, of which meaningfulness is debatable for semantic characterization [3]. We therefore make use of grammatical features, for partial semantic characterization of object affordances. While other semantic relationships employed in engineering are not easily prone to confident, automatic extraction and knowledge engineers have to recur to manual ontology insertion [4], the author's claim is that potentiality relationships can be robustly extracted from grammar relationships of Subject-Verb-Object (SVO) co-occurrences. The choice of the ability modality relationship calls for the assumption that the training corpus from which we derive data has to have reliable, non-figurative subject-verb-object co-occurrence tuples. More formally, co-occurrence of every noun $s_1 \in S$ with a verb $v_1 \in V$ entails the ability of s_1 to perform such action v_1 on the co-occurring object $o_1 \in O$. In simpler terms, we assume the instance "a robot builds a desk" implies "*a robot can build*" and "*a desk is buildable*", which the author does not consider a restrictive assumption that requires controlled authoring.

[1] Computer Science Department, Technische Universität München, Germany nicholas.kirk@tum.de

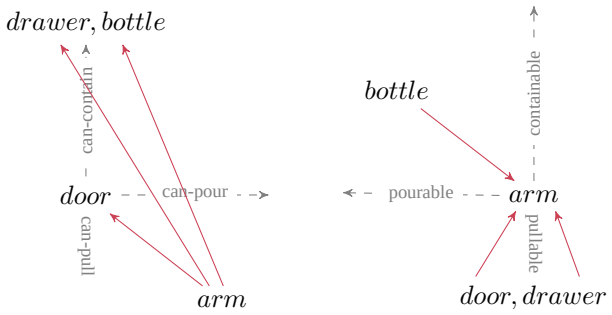

Fig. 1: Hypothetical representation of a dual active/passive 3-dimensional modality space (with predicates as dimensions) representing instances of kitchen scenario objects.

We therefore model our symbolic knowledge on potentiality as the joint probability distribution of all SVO occurrences in our training source, obtained via learning on typed dependency analysis output features of such source (see Eq. 1).

$$Modality(W) = P\left(S \times V \times O\right) \tag{1}$$

$$S = \{\forall s \in N \mid grammar_type(s, subject)\}$$
$$V = \{\forall v \in N \mid grammar_type(v, verb)\}$$
$$O = \{\forall o \in N \mid grammar_type(o, object)\}$$

From Eq. 1 we derive two dual joint probability distributions, which encapsulate knowledge of active and passive noun roles (Eq. 2 and Eq. 3) and can induce two distinct vector spaces, representing passive and active role information (example in Fig. 1).

$$Modality_{\text{active}}(W) = P(S \times V) \tag{2}$$
$$Modality_{\text{passive}}(W) = P(V \times O) \tag{3}$$

II. Implementation

In order to learn our distribution in Eq. 1, a possible approach is to exploit Markov Logic Networks (MLN) [5] on a set of previously extracted Stanford typed dependencies [6]. The latter are a labeled, directed grammar relationship among pairs of words, which capture word order and relationship type (Fig. 2): when considering `'nsubj'` (subject of an action) and `'dobj'` (object of an action) labels, these can be seen as grounded action-object predicates.

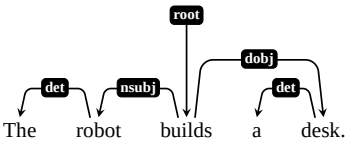

Fig. 2: Example of words that compose a sentence instance and their typed dependencies (illustrated as labeled directed edges).

We can then perform learning on such grounded models thanks to MLN, which is a knowledge representation formalism that enables probabilistic learning and inference via the combined use of first order logic and probabilistic undirected graphical models (i.e. *Markov Random Fields*). More formally, MLN theory defines a probability over the world x as a log-linear model in which we have an exponentiated sum of weights w_j of a binary feature f_j, and the partition function Z (see Eq. 4).

$$P\left(X = x\right) = \frac{1}{Z}\ exp\left(\sum_{j} w_j f_j\left(x\right)\right) \tag{4}$$

In our case, we consider the binary formula $f_j(x)$ as an evaluation of a logic formula representing grammar relations as predicates, and we substitute such term with $n_j(x)$, where the latter is number of true groundings of such formula f_j in x_j. The MLN formalism aims to learn the stationary distribution (i.e. learn stationary weight values w_j) of the true groundings $n_j(x)$, possibly a sufficient heuristic condition for scalability.

III. Discussion

a) Related Work: Systems which focus on the initialization parameters from ontologies (i.e. aggregates of semantic relationships and entities) do not debate how such source was populated [7]. Some previous literature does value mappings between language constructs and affordances, but analyze the opposite problem [8]. Closer work which adopts MLN and grammar features has been proven successful for mining natural language instructions for the robotics domain [9], but does not focus on affordance understanding and concentrates on inferring likely action roles, while other literature does make use of MLN, but does not employ grammar feature analysis [10]. Closer work does consider typed dependency extraction for semantic characterization, but does not focus on SVO tuple analysis [11], [12].

b) Evaluation: As the system can process a high number of noun-action relationships, we require an equally well populated ontology representing ground truth references. For activity and passivity labels, the scope might require manual annotation.

c) Potential: Other than fulfilling the requirement of providing an initial affordance world estimate, it can provide understanding of hidden or partially observable affordances [13], particularly useful when objects are not in full reach of the perception array. The vector space induction enables the use high-dimensional tensor computations for semantic characterization adopted in linguistics (such as compositionality and retrieval of neighboring entries [14]), to a yet unknown extent of effectiveness within the context. The accuracy of the learnt ontology entries can be refined via reinforcement learning during the course of task executions, or be verified thanks to an active learning process. More precisely, the SVO entries can be encapsulated in questions in a process called language generation, in order to ask for their correctness [15].

d) Limitations: Although we assume the text is confined to a technical domain, the authors of the source might make use of partly figurative wordings. As a result, the word frequency distribution would present bias or outliers (i.e. presence of erroneous co-occurrences of analyzed nouns or figurative nouns unrelated to known entities). Furthermore, also the independent word frequency of occurrence does not provide information regarding entity existence, and would require a form of normalization.

e) Conclusions: The linguistic and computational obstacles towards model effectiveness are manifold, and surely require the development of processes such as bias removal and outlier detection. However, this concept paper highlights the usage of technical text mining for affordance acquisition, and mainly points to the potential of induced vector spaces for retrieving objects with similar affordance, or the affordance of aggregates, and above all its practical use as initial world affordance estimate.

References

[1] T. E. Horton, A. Chakraborty, and R. St. Amant, "Affordances for robots: a brief survey," *AVANT. Pismo Awangardy Filozoficzno-Naukowej*, no. 2, pp. 70–84, 2012.

[2] Z. S. Harris, "Distributional structure." *Word*, 1954.

[3] M. Sahlgren, "The distributional hypothesis," *Italian Journal of Linguistics*, vol. 20, no. 1, pp. 33–54, 2008.

[4] L. Reeve and H. Han, "Survey of semantic annotation platforms," in *Proceedings of the 2005 ACM symposium on Applied computing*. ACM, 2005, pp. 1634–1638.

[5] M. Richardson and P. Domingos, "Markov logic networks," *Machine learning*, vol. 62, no. 1-2, pp. 107–136, 2006.

[6] M.-C. de Marneffe and C. D. Manning, "The stanford typed dependencies representation," in *22nd International Conference on Computational Linguistics*, 2008, p. 1.

[7] S. S. Hidayat, B. K. Kim, and K. Ohba, "Learning affordance for semantic robots using ontology approach," in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 2008, pp. 2630–2636.

[8] O. Yürüten, K. F. Uyanık, Y. Çalışkan, A. K. Bozcuoğlu, E. Şahin, and S. Kalkan, "Learning adjectives and nouns from affordances on the icub humanoid robot," in *From Animals to Animats 12*. Springer, 2012, pp. 330–340.

[9] D. Nyga and M. Beetz, "Everything robots always wanted to know about housework (but were afraid to ask)," in *Intelligent Robots and Systems (IROS), 2012 IEEE/RSJ International Conference on*. IEEE, 2012, pp. 243–250.

[10] I. Beltagy, C. Chau, G. Boleda, D. Garrette, K. Erk, and R. Mooney, "Montague meets markov: Deep semantics with probabilistic logical form," in *2nd Joint Conference on Lexical and Computational Semantics: Proceeding of the Main Conference and the Shared Task, Atlanta*, 2013, pp. 11–21.

[11] S. Padó and M. Lapata, "Dependency-based construction of semantic space models," *Computational Linguistics*, vol. 33, no. 2, pp. 161–199, 2007.

[12] G. Boella, L. Di Caro, and L. Robaldo, "Semantic relation extraction from legislative text using generalized syntactic dependencies and support vector machines," *Theory, Practice, and Applications of Rules on the Web*, p. 218.

[13] W. W. Gaver, "Technology affordances," in *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 1991, pp. 79–84.

[14] T. Mikolov, K. Chen, G. Corrado, and J. Dean, "Efficient estimation of word representations in vector space," *arXiv preprint arXiv:1301.3781*, 2013.

[15] N. H. Kirk, D. Nyga, and M. Beetz, "Controlled natural languages for language generation in artificial cognition," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE RAS, 2014, pp. 6667–6672.